arXiv:1602.01636v2 [math.NA] 21 Jul 2016

ISOGEOMETRIC PRECONDITIONERS BASED ON FAST SOLVERS FOR THE SYLVESTER EQUATION*

GIANCARLO SANGALLI† AND MATTIA TANI†

Abstract. We consider large linear systems arising from the isogeometric discretization of the Poisson problem on a single-patch domain. The numerical solution of such systems is considered a challenging task, particularly when the degree of the splines employed as basis functions is high. We consider a preconditioning strategy which is based on the solution of a Sylvester-like equation at each step of an iterative solver. We show that this strategy, which fully exploits the tensor structure that underlies isogeometric problems, is robust with respect to both mesh size and spline degree, although it may suffer from the presence of complicated geometry or coefficients. We consider two popular solvers for the Sylvester equation, a direct one and an iterative one, and we discuss in detail their implementation and efficiency for 2D and 3D problems on single-patch or conforming multi-patch NURBS geometries.

Numerical experiments for problems with different domain geometries are presented, which demonstrate the potential of this approach.

Key words. Isogeometric analysis, preconditioning, Kronecker product, Sylvester equation.

AMS subject classifications. 65N30, 65F08, 65F30

1. Introduction.

The Isogeometric method is a computational technique for solving partial differential equations (PDEs). It has been proposed in the seminal paper [27] as an extension of the classical finite element method, and is based on the idea of using splines or other functions constructed from splines (e.g., non-uniform rational B-splines, NURBS) both for the parametrization of the computational domain, as it is typically done by computer aided design software, and for the representation of the unknown solution fields of the PDE of interest. Many papers have demonstrated the effective advantage of isogeometric methods in various frameworks. The interested reader may find a detailed presentation of this idea with engineering applications in the book [13].

Unlike standard finite element methods, the isogeometric method makes it possible to use high-regularity functions. The so-called isogeometric k-method, based on splines of degree p and global C^{p-1} regularity, has shown significant advantages in term of higher accuracy per degree-of-freedom in comparison to C^0 finite elements of degree p [14, 4]. However, the computational cost per degree-of-freedom is also higher for the k-method, in currently available isogeometric codes. In practice, quadratic or cubic splines are typically preferred as they maximise computational efficiency.

Standard isogeometric codes typically re-use finite element technology, which is very convenient but at the same time not the best choice for computational efficiency. The two fundamental stages of a linear PDE solver are the formation of the system matrix $\mathcal{A}$ and the solution of the linear system $\mathcal{A}u = b$, and both stages, in standard isogeometric software, show a computational cost that grows significantly as the degree p grows. The focus of this paper is on the second stage, that is the linear solver.

The study of the computational efficiency of linear solvers for isogeometric discretizations has been initiated in the papers [12, 11], where it has been shown that the algorithms used with the finite element method suffer of performance degradation when used to solve isogeometric linear systems. Consider, for example, a Lagrangian

*Version of February 22, 2016

†Università di Pavia, Dipartimento di Matematica "F. Casorati", Via A. Ferrata 1, 27100 Pavia, Italy. Emails: {sangia05,mattia.tani}@unipv.it

finite element method with polynomial degree p and N degrees-of-freedom, in 3D: the system matrix $\mathcal{A}$ has a storage cost of $O(Np^3)$ non-zero terms and a solving cost by a direct solver of $O(N^2)$ floating point operations (FLOPs), (see [12, Section 2.3], under the assumption $N > p^6$). If, instead, we consider the isogeometric k-method with C^{p-1} p-degree splines and N degrees-of-freedom, the system matrix $\mathcal{A}$ has still $O(Np^3)$ non-zero entries, but a standard direct solver costs $O(N^2p^3)$ FLOPs, i.e., p^3 times more than a finite element approximation.

Iterative solvers have attracted more attention in the isogeometric community. The effort has been primarily on the development of preconditioners for the Poisson model problem, for arbitrary degree and continuity splines. As reported in [11], standard algebraic preconditioners (Jacobi, SSOR, incomplete factorization) commonly adopted for finite elements exhibit reduced performance when used in the context of the isogeometric k-method. Multilevel and multigrid approaches are studied respectively in [10] and [22], while advances in the theory of domain-decomposition based solvers are given in, e.g., [5, 6, 9]. These papers also confirm the difficulty in achieving both *robustness* and computational *efficiency* for the high-degree k-method. In this context, we say that a preconditioner $\mathcal{P}$ for the linear system $\mathcal{A}u = b$ is robust if the condition number $\kappa\left(\mathcal{P}^{-1}\mathcal{A}\right)$ is bounded from above by a reasonably low number, independent of the degree or continuity of the spline space adopted in the isogeometric discretization; we say that a preconditioner is computationally efficient if its setup and application has a computational cost comparable to the one of the matrix-vector product for the system matrix $\mathcal{A}$, i.e. $O(Np^3)$.

More sophisticated multigrid preconditioners have been proposed in the recent papers [19] and [26]. The latter, in particular, contains a proof of robustness, based on the theory of [38]. The two works ground on the following common ingredients: specific spectral properties of the discrete operator of the isogeometric k-method and the tensor-product structure of isogeometric spaces.

In this paper we also exploit the tensor-product structure of multivariate spline space, on a different basis. We rely on approaches that have been developed for the so-called Sylvester equation.

Consider the Laplace operator with constant coefficients, on the square $[0,1]^2$, then the tensor-product spline Galerkin discretization leads to the system

$$(K_1 \otimes M_2 + M_1 \otimes K_2)u = b, \tag{1.1}$$

where K_ℓ and M_ℓ denote the univariate stiffness and mass matrices in the ℓ direction, $\ell = 1, 2$, and $\otimes$ is the Kronecker product. Equation (1.1), when reformulated as a matrix equation, takes the name of (generalized) Sylvester equation. This is a well studied problem in the numerical linear algebra literature, as it appears in many applications, e.g. stochastic PDEs, control theory, etc. (see the recent survey [36]). Observe that in general, for variable coefficients, general elliptic problems, non trivial and possibly multi-patch geometry parametrization, the isogeometric system is not as in (1.1). In this case, a fast solver for (1.1) plays the role of a preconditioner. Having this motivation in mind, our aim is to study how the linear solvers for the Sylvester equation perform for (1.1), especially when originated by an isogeometric k-method.

We select two among the most popular algorithms: the first is the fast diagonalization (FD) direct solver proposed by Lynch, Rice and Thomas in [33], the second is the alternating direction implicit (ADI) iterative solver, first introduced in [34] and further developed in a number of papers, among which [40]. The potential of ADI in the context of isogeometric problems has already been recognized in [23].

Our ultimate goal is the solution of 3D isogeometric systems, especially when high resolution is needed. A remarkable example is the simulation of turbulence, see e.g. [3]. Here the Poisson problem on a unit cube leads to the linear system

$$(K_1 \otimes M_2 \otimes M_3 + M_1 \otimes K_2 \otimes M_3 + M_1 \otimes M_2 \otimes K_3)u = b,$$

which can be efficiently solved with the FD method and a generalization of ADI.

We analyze and benchmark the proposed approaches in both 2D and 3D cases. The results show that the FD method exceeds by far ADI in terms of computational efficiency. This is seen in 2D, but the gap is much wider in 3D. Here, the FD solver count is $O(N^{4/3})$ FLOPs. ADI costs $O(Np)$ FLOPs per iteration, which result in an asymptotically lower operation counting. However, when used as a preconditioner in benchmarks that are representative of realistic problems, the FD solver performs orders of magnitude better. In fact, in all our benchmarks that uses a conjugate gradient (CG) iterative solver, the computational time spent in the FD preconditioner application is even lower than the residual computation (multiplication of matrix $\mathcal{A}$ times a vector). This surprising performance is due to the fact that the FD solver requires dense matrix-matrix operations that takes advantages of modern computer architecture. In particular, the performance boost is due to the efficient use of the CPU hierarchy cache and memory access. Furthermore, the FD method is especially suited to parallelization which may significantly speed up the execution time, though this is not considered in our analysis.

Concerning robustness, in both approaches the condition number $\kappa\left(\mathcal{P}^{-1}\mathcal{A}\right)$ depends only on the geometry parametrization of the computational domain Ω and on the coefficients of the differential operator. In this paper we study this dependence and perform some numerical tests. We will show that a singular mapping causes a loss of robustness, while $\kappa\left(\mathcal{P}^{-1}\mathcal{A}\right)$ is uniformly bounded w.r.t. the degree p and mesh size h when the parametrization is regular. In all cases, it is important to have strategies to further improve the condition number, but this goes beyond the scope of the present work and this will be the topic of further researches.

Finally, we show how to combine the considered preconditioners with a domain decomposition approach in order to solve multi-patch problems with conforming discretization. The overall strategy naturally inherits efficiency and robustness from the preconditioners for single-patch problems.

The outline of the paper is as follows. In Section 2 we introduce the matrices stemming from isogeometric discretization; we also recall the Kronecker product notation and its main properties. In Section 3 we define the preconditioner and discuss the spectral condition number of the preconditioned system. Sections 4 and 5 describe how such preconditioner can be efficiently applied in the 2D and 3D case, respectively. In Section 6, the multi-patch setting is discussed. Numerical experiments are reported in Section 7. Finally, in Section 8 we draw the conclusions and outline future research directions.

2. Preliminaries.

2.1. Splines-based isogeometric method. We consider, as a model problem, the Poisson equation with Dirichlet boundary conditions:

$$\begin{cases} -\mathrm{div}(K(\mathbf{x})\nabla \mathrm{u}(\mathbf{x})) = f(\mathbf{x}) & \text{on} \quad \Omega, \\ \mathrm{u} = 0 & \text{on} \quad \partial\Omega, \end{cases} \tag{2.1}$$

where $K(\mathbf{x})$ is a symmetric positive definite matrix for each $\mathbf{x} \in \Omega$. In isogeometric methods, Ω is given by a spline or NURBS parametrization. For the sake of simplicity, we consider a single-patch spline parametrization.

Given two positive integers p and m, we say that $\Xi := \{\xi_1, \dots, \xi_{m+p+1}\}$ is an open knot vector if

$$\xi_1 = \dots = \xi_{p+1} < \xi_{p+2} \leq \dots \leq \xi_m < \xi_{m+1} = \dots = \xi_{m+p+1}, \tag{2.2}$$

where repeated knots are allowed, up to multiplicity p. Without loss of generality, we assume $\xi_1 = 0$ and $\xi_{m+p+1} = 1$. From the knot vector Ξ, B-spline functions of degree p are defined e.g. by the Cox-DeBoor recursive formula: piecewise constants ($p = 0$) B-splines are

$$\hat{B}_{i,0}(\zeta) = \begin{cases} 1 & \text{if } \xi_i \leq \zeta < \xi_{i+1}, \\ 0 & \text{otherwise,} \end{cases} \tag{2.3}$$

and for $p \geq 1$ the *B-spline* functions are obtained by the recursion

$$\hat{B}_{i,p}(\zeta) = \frac{\zeta - \xi_i}{\xi_{i+p} - \xi_i} \hat{B}_{i,p-1}(\zeta) + \frac{\xi_{i+p+1} - \zeta}{\xi_{i+p+1} - \xi_{i+1}} \hat{B}_{i+1,p-1}(\zeta), \tag{2.4}$$

where $0/0 = 0$. In general, the B-spline functions are degree p piecewise polynomial with $p - r$ continuous derivative at each knot with multiplicity r. In this work we are primarily interested in the so called k-refinement or isogeometric k-method, see [13]. For that, we assume that the multiplicity of all internal knots is 1, which corresponds to C^{p-1} continuous splines. Each B-spline $\hat{B}_{i,p}$ depends only on $p + 2$ knots, which are collected in the *local knot vector*

$$\Xi_{i,p} := \{\xi_i, \dots, \xi_{i+p+1}\}.$$

When needed, we will adopt the notation $\hat{B}_{i,p}(\zeta) = \hat{B}[\Xi_{i,p}](\zeta)$. The support of each basis function is exactly $\text{supp}(\hat{B}_{i,p}) = [\xi_i, \xi_{i+p+1}]$.

Multivariate B-splines in dimension d ($d = 2, 3$ are the interesting cases) are defined from univariate B-splines by tensorization. For the sake of simplicity, we assume that the degree p and the length of the knot vectors $\Xi^\ell = \{\xi_{\ell,1}, \dots, \xi_{\ell,m+p+1}\}$, is the same in all directions $\ell = 1, \dots, d$. Then for each multi-index $\boldsymbol{i} = (i_1, \dots, i_d)$, we introduce the local knot vector $\Xi^1_{i_1,p} \times \dots \times \Xi^d_{i_d,p}$ and the multivariate B-spline

$$\hat{B}_{\boldsymbol{i},p}(\boldsymbol{\zeta}) = \hat{B}[\Xi^1_{i_1,p}](\zeta_1) \dots \hat{B}[\Xi^d_{i_d,p}](\zeta_d). \tag{2.5}$$

To simplify the notation, when not needed the subscript p is not indicated.

The domain Ω is given by a d-dimensional single-patch spline parametrization

$$\Omega = \boldsymbol{F}([0,1]^d), \text{ with } \boldsymbol{F}(\boldsymbol{\zeta}) = \sum_{\boldsymbol{i}} \boldsymbol{C_i} \hat{B}_{\boldsymbol{i}}(\boldsymbol{\zeta}),$$

where $\boldsymbol{C_i}$ are the control points. Following the isoparametric paradigm, the basis functions $B_{\boldsymbol{i}}$ on Ω are defined as $B_{\boldsymbol{i}} = \hat{B}_{\boldsymbol{i}} \circ \boldsymbol{F}^{-1}$. The isogeometric space, incorporating the homogeneous Dirichlet boundary condition, reads

$$V_h = \text{span}\{B_{\boldsymbol{i}} \text{ such that } \boldsymbol{i} = (i_1, \dots, i_d), \text{ with } 2 \leq i_\ell \leq m - 1, 1 \leq \ell \leq d\}. \tag{2.6}$$

We introduce a scalar indexing for functions in (2.6) as follows: to the each multi-index $\boldsymbol{i} = (i_1, \dots, i_d)$ we associate $i = 1 + \sum_{\ell=1}^{d} (m-2)^{\ell-1}(i_\ell - 2)$ and, with abuse of

notation, indicate $B_{\boldsymbol{i}} = B_i$, etc. The dimension of V_h is denoted as $N = n^d$, where $n = m - 2$. Then, the Galerkin stiffness matrix reads

$$\begin{aligned}\mathcal{A}_{ij} &= \int_\Omega \left(\nabla B_i(\boldsymbol{x})\right)^T K(\boldsymbol{x}) \nabla B_j(\boldsymbol{x}) d\boldsymbol{x} \\ &= \int_{[0,1]^d} \left(\nabla \hat{B}_i(\boldsymbol{\xi})\right)^T Q(\boldsymbol{\xi})\, \nabla \hat{B}_j\left(\boldsymbol{\xi}\right)\, d\boldsymbol{\xi}, \qquad i,j = 1,\dots,N\end{aligned} \tag{2.7}$$

where

$$Q = \det\left(J_{\boldsymbol{F}}\right) J_{\boldsymbol{F}}^{-T} K J_{\boldsymbol{F}}^{-1} \tag{2.8}$$

and $J_{\boldsymbol{F}}$ denotes the Jacobian of $\boldsymbol{F}$.

The support of a B-spline in V_h that does not touch $\partial\Omega$ intersects the support of $(2p+1)^d$ splines in V_h (including itself). If the support of a B-splines intersects with $\partial\Omega$, it overlaps at least $(p+1)^d$ and up to $(2p+1)^d$ B-spline supports (including itself). Thus, the number of nonzeros of $\mathcal{A}$ is about $(2p+1)^d N$.

2.2. Kronecker product. Let $A \in \mathbb{R}^{n_a \times n_a}$, and $B \in \mathbb{R}^{n_b \times n_b}$. The *Kronecker product* between A and B is defined as

$$A \otimes B = \begin{bmatrix} a_{11}B & \dots & a_{1n_a}B \\ \vdots & \ddots & \vdots \\ a_{n_a 1}B & \dots & a_{n_a n_a}B \end{bmatrix} \in \mathbb{R}^{n_a n_b \times n_a n_b},$$

where a_{ij}, $i,j = 1,\dots n_a$, denote the entries of A. The Kronecker product is an associative operation, and it is bilinear with respect to matrix sum and scalar multiplication. We now list a few properties of the Kronecker product that will be useful in the following.

- It holds

$$(A \otimes B)^T = A^T \otimes B^T. \tag{2.9}$$

- If C and D are matrices of conforming order, then

$$(A \otimes B)\,(C \otimes D) = (AC \otimes BD). \tag{2.10}$$

- If A and B are nonsingular, then

$$(A \otimes B)^{-1} = A^{-1} \otimes B^{-1}. \tag{2.11}$$

- If λ_i, $i = 1,\dots,n_a$, denote the eigenvalues of A and μ_b, $j = 1,\dots,n_2$, denote the eigenvalues of B, then the $n_a n_b$ eigenvalues of $A \otimes B$ have the form

$$\lambda_i \mu_j, \qquad i = 1,\dots,n_a,\ j = 1,\dots,n_b. \tag{2.12}$$

Property (2.9) implies that if A and B are both symmetric, then $A \otimes B$ is also symmetric. Moreover, if A and B are both positive definite, then according to (2.12) $A \otimes B$ is also positive definite.

For any matrix $X \in \mathbb{R}^{n_a \times n_b}$ we denote with $\mathrm{vec}(X)$ the vector of $\mathbb{R}^{n_a n_b}$ obtained by "stacking" the columns of X. Then if A, B and X are matrices of conforming order, and $x = \mathrm{vec}(X)$, it holds

$$(A \otimes B)\, x = \mathrm{vec}(BXA^T). \tag{2.13}$$

This property can be used to cheaply compute matrix-vector products with a matrix having Kronecker structure. Indeed, it shows that computing $(A \otimes B)\, x$ is equivalent to computing n_b matrix-vector products with A and n_a matrix-vector products with B. Note in particular that $(A \otimes B)$ does not have to be formed.

If A and B are nonsingular, then (2.13) is equivalent to

$$(A \otimes B)^{-1}\, x = \mathrm{vec}(A^{-1} X B^{-T}), \tag{2.14}$$

which, in a similar way, shows that the problem of solving a linear system having $(A \otimes B)$ as coefficient matrix is equivalent to solve n_b linear systems involving A and n_a linear systems involving B.

2.3. Evaluation of the computational cost and efficiency. Throughout the paper, we will primarily evaluate the computational cost of an algorithm by counting the number of floating point operations (FLOPs) it requires. A single addition, subtraction, multiplication or division performed in floating point arithmetic counts as one FLOP [24]. The number of FLOPs associated with an algorithm is an indication to assess its efficiency, and it is widely employed in literature. However, any comparison of FLOPs between different algorithms should be interpreted with caution. We emphasize, indeed, that the number of FLOPs represents just a portion of the computational effort required by an algorithm, as it does not take into account the movement of data in the memory and other overheads that affect the execution time. While these are difficult to estimate, we will discuss them when needed.

3. The preconditioner. Consider the matrix

$$\mathcal{P}_{ij} = \int_{[0,1]^d} \left(\nabla \hat{B}_i\right)^T \nabla \hat{B}_j \, d\xi, \qquad i,j = 1,\dots,N. \tag{3.1}$$

Observe that $\mathcal{P} = \mathcal{A}$ in the special case when K and $J_{\boldsymbol{F}}$ are the identity matrices, which means in particular that Ω is the unit cube. For $d = 2$, by exploiting the tensor product structure of the basis functions we have

$$\mathcal{P} = K_1 \otimes M_2 + M_1 \otimes K_2,$$

where M_1, M_2 represent the mass, and K_1, K_2 the stiffness univariate matrices.

$$(M_1)_{ij} = \int_0^1 \hat{B}[\Xi_i^1](\zeta_1)\, \hat{B}[\Xi_j^1](\zeta_1)\, d\zeta_1, \qquad (M_2)_{ij} = \int_0^1 \hat{B}[\Xi_i^2](\zeta_2)\, \hat{B}[\Xi_j^2](\zeta_2)\, d\zeta_2,$$

$$(K_1)_{ij} = \int_0^1 (\hat{B}[\Xi_i^1])'(\zeta_1) \cdot (\hat{B}[\Xi_j^1])'(\zeta_1)\, d\zeta_1,$$

$$(K_2)_{ij} = \int_0^1 (\hat{B}[\Xi_i^2])'(\zeta_2) \cdot (\hat{B}[\Xi_j^2])'(\zeta_2)\, d\zeta_2.$$

Such matrices are all symmetric positive definite and banded with bandwidth p (we say that a matrix B has bandwidth p if $B_{ij} = 0$ for $|i-j| > p$). These matrices have the same order $n = m - 2$, where m is the dimension of the univariate spline space (see also from (2.2)). Similarly, when $d = 3$

$$\mathcal{P} = K_1 \otimes M_2 \otimes M_3 + M_1 \otimes K_2 \otimes M_3 + M_1 \otimes M_2 \otimes K_3.$$

By comparing (2.7) and (3.1), observe that $\mathcal{P}_{ij} \neq 0$ if and only if $\mathcal{A}_{ij} \neq 0$. Thus, despite having different entries in general, $\mathcal{A}$ and $\mathcal{P}$ have the same sparsity pattern.

We propose $\mathcal{P}$, defined in (3.1), as a preconditioner for the isogeometric matrix $\mathcal{A}$. In other words, we want to precondition a problem with arbitrary geometry and coefficients with a solver for the same operator on the parameter domain, with constant coefficients. This is a common approach, see e.g. [23], [19] and [26].

Note that, according to (2.9) and (2.12), $\mathcal{P}$ is symmetric and positive definite, and hence we can use it as preconditioner for the conjugate gradient (CG) method. At each CG iteration, we need to solve a system of the form

$$\mathcal{P}s = r, \tag{3.2}$$

where r is the current residual. Due to the structure of $\mathcal{P}$, (3.2) is a Sylvester-like equation. How to efficiently solve this system for $d = 2$ and $d = 3$, employing solvers for Sylvester equation, will be the topic of Sections 4 and 5. In this section we discuss the effects of geometry and coefficients on the overall CG convergence. The next proposition provides an upper bound for the spectral condition number of $\mathcal{P}^{-1}\mathcal{A}$.

PROPOSITION 1.

It holds

$$\kappa\left(\mathcal{P}^{-1}\mathcal{A}\right) \leq \frac{\sup\limits_{\Omega} \lambda_{\max}(Q)}{\inf\limits_{\Omega} \lambda_{\min}(Q)}, \tag{3.3}$$

where the matrix Q is given in (2.8).

Proof.

Let $u = (u_1, \ldots, u_N)^T \in \mathbb{R}^n$, and define $u_h = \sum_{i=1}^{N} u_i \hat{B}_i$. Then it holds

$$u^T \mathcal{A} u = \int_{[0,1]^d} (\nabla u_h)^T Q \, \nabla u_h \, d\xi \leq \int_{[0,1]^d} \lambda_{\max}(Q) \, \|\nabla u_h\|^2 \, d\xi$$

$$\leq \sup_{\Omega} \lambda_{\max}(Q) \int_{[0,1]^d} \|\nabla u_h\|^2 \, d\xi = \sup_{\Omega} \lambda_{\max}(Q) \; u^T \mathcal{P} u.$$

By the Courant-Fischer theorem, we infer $\lambda_{\max}\left(\mathcal{P}^{-1}\mathcal{A}\right) \leq \sup\limits_{\Omega} \lambda_{\max}(Q)$. With analogous calculations one can show that $\lambda_{\min}\left(\mathcal{P}^{-1}\mathcal{A}\right) \geq \inf\limits_{\Omega} \lambda_{\min}(Q(\xi))$, and hence $\kappa\left(\mathcal{P}^{-1}\mathcal{A}\right) \leq \frac{\sup\limits_{\Omega} \lambda_{\max}(Q)}{\inf\limits_{\Omega} \lambda_{\min}(Q)}$. □

Proposition 1 states a useful and well-known result, that formalizes an intuitive fact. As long as the considered problem does not depart much from the model problem on the square with constant coefficients, the right-hand side of (3.3) will be small and the preconditioner is expected to perform well. On the other hand, if the eigenvalues of Q vary widely, due to the presence of complicated geometry or coefficients, the preconditioner performance decreases. In these cases, it is useful to have strategies to improve the spectral conditioning of $\mathcal{P}^{-1}\mathcal{A}$: this is a topic that we will address in a forthcoming paper. We emphasize that bound (3.3) does not depend neither on the mesh size nor on the spline degree, but only on $\boldsymbol{F}$ and K.

Furthermore Proposition 1 allows one to compare the strategies proposed in this paper with different approaches which do not rely on the preconditioner $\mathcal{P}$. This might

be helpful not only from the theoretical point of view, but also from a practical perspective. Indeed, during the process of assembling the stiffness matrix $\mathcal{A}$, the matrix Q can be evaluated at all quadrature points of the mesh and the extreme eigenvalues of Q can be computed in order to estimate the right-hand side of (3.3). This leads to a reliable estimate of $\kappa\left(\mathcal{P}^{-1}\mathcal{A}\right)$ before attempting to solve the system. If another solver is available, which does not suffer from complicated geometry or coefficients, then a smart software could use the estimate on $\kappa\left(\mathcal{P}^{-1}\mathcal{A}\right)$ to automatically choose which method is more suited to solve the system at hand.

4. The 2D case. When $d = 2$, equation (3.2) takes the form

$$(K_1 \otimes M_2 + M_1 \otimes K_2)\, s = r. \tag{4.1}$$

Using relation (2.13), we can rewrite this equation in matrix form

$$M_2 S K_1 + K_2 S M_1 = R, \tag{4.2}$$

where $\mathrm{vec}(S) = s$ and $\mathrm{vec}(R) = r$. Equation (4.2) takes the name of (generalized) Sylvester equation. Due to its many applications, the literature dealing with Sylvester equation (and its variants) is vast, and a number of methods have been proposed for its numerical solution. We refer to [36] for a recent survey on this subject.

In the last two decades, the research on Sylvester equation has mainly focused on methods which require that the right-hand side matrix R has low rank. Such methods are nor considered in this work. Indeed, even if there are cases where R is low-rank or can be approximated efficiently by a low-rank matrix, this is not the general case. Furthermore, and perhaps more important, low-rank methods are designed for solving very large problems, where even storing the solution S might be unfeasible. This is not the case of problems of practical interest in isogeometric analysis.

In this paper, we consider two among the most studied methods, which in the authors' perspective seem the most suited for the particular features of IGA problems. The fast diagonalization (FD) method is a direct solver, which means that $s = \mathcal{P}^{-1}r$ is computed exactly. The alternating direction implicit (ADI) method is an iterative solver, which means that s is computed only approximately. We remark that ADI was first applied to IGA problems in [23].

To keep the notation consistent with the rest of the paper, in this section we will favor the Kronecker formulation (4.1) with respect to the matrix equation form (4.2).

4.1. The 2D fast diagonalization method. We describe a direct method for (4.1) that was first presented in 1964 by Lynch, Rice and Thomas [33] as a method for solving elliptic partial differential equations discretized with finite differences. Following [18], we refer to it as the fast diagonalization (FD) method. We remark that this approach was extended to a general Sylvester equation involving nonsymmetric matrices by Bartels and Stewart in 1972 [2], although this is not considered here.

We consider the generalized eigendecomposition of the matrix pencils (K_1, M_1) and (K_2, M_2), namely

$$K_1 U_1 = M_1 U_1 D_1 \qquad K_2 U_2 = M_2 U_2 D_2, \tag{4.3}$$

where D_1 and D_2 are diagonal matrices whose entries are the eigenvalues of $M_1^{-1}K_1$ and $M_2^{-1}K_2$, respectively, while U_1 and U_2 satisfy

$$U_1^T M_1 U_1 = I, \qquad U_2^T M_2 U_2 = I,$$

which implies in particular $U_1^{-T}U_1^{-1} = M_1$ and $U_2^{-T}U_2^{-1} = M_2$, and also, from (4.3), $U_1^{-T}D_1U_1^{-1} = K_1$ and $U_2^{-T}D_2U_2^{-1} = K_2$. Therefore we factorize $\mathcal{P}$ in (4.1) as follows:

$$(U_1 \otimes U_2)^{-T} (D_1 \otimes I + I \otimes D_2) (U_1 \otimes U_2)^{-1} s = r,$$

and adopt the following strategy:

Algorithm 1 FD direct method (2D)

1: Compute the generalized eigendecompositions (4.3)
2: Compute $\tilde{r} = (U_1 \otimes U_2)^T r$
3: Compute $\tilde{s} = (D_1 \otimes I + I \otimes D_2)^{-1} \tilde{r}$
4: Compute $s = (U_1 \otimes U_2)\tilde{s}$

Computational cost. The exact cost of the eigendecompositions in line 1 depends on the algorithm employed. We refer to [24, Chapter 8], [17, Section 5.3] and references therein for an overview of the state-of-the-art methods. A simple approach is to first compute the Cholesky factorization $M_1 = LL^T$ and the symmetric matrix $\widetilde{K_1} = L^{-1}K_1L^{-T}$. Since M_1 and K_1 are banded, the cost of these computations is $\mathrm{O}(pn^2)$ FLOPs. The eigenvalues of $\widetilde{K_1}$ are the same of (4.3), and once the matrix $\widetilde{U_1}$ of orthonormal eigenvectors is computed then one can compute $U_1 = L^{-T}\widetilde{U_1}$, again at the cost of $\mathrm{O}(pn^2)$ FLOPs. Being $\widetilde{U_1}$ orthogonal, then $U_1^T M_1 U_1 = I_n$. If the eigendecomposition of $\widetilde{K_1}$ is computed using a divide-and-conquer method, the cost of this operation is roughly $4n^3$ FLOPs. We remark that the divide-and-conquer approach is also very suited for parallelization. In conclusion, by this approach, line 1 requires roughly $8n^3$ FLOPs.

Lines 2 and 4 each involve a matrix-vector product with a matrix having Kronecker structure, and each step is equivalent (see (2.13)) to 2 matrix products involving dense $n \times n$ matrices. The total computational cost of both steps is $8n^3$ FLOPs. Line 3 is just a diagonal scaling, and its $O(n^2)$ cost is negligible. We emphasize that the overall computational cost of Algorithm 1 is independent of p.

If we apply Algorithm 1 as a preconditioner, then Step 1 may be performed only once, since the matrices involved do not change throughout the CG iteration. In this case the main cost can be quantified in approximately $8n^3$ FLOPs per CG iteration. The other main computational effort of each CG iteration is the residual computation, that is the product of the system matrix $\mathcal{A}$ by a vector, whose cost in FLOPs is twice the number of nonzero entries of $\mathcal{A}$, that is approximately $2(2p+1)^2n^2$. In conclusion, the cost ratio between the preconditioner application and the residual computation is about $4n/(2p+1)^2 \approx n/p^2$.

4.2. The ADI method. The ADI method was originally proposed in 1955 by Peaceman and Racheford as a method to solve elliptic and parabolic differential equations with two space variables [34] by a structured finite difference discretization. An important contribution to the early development of ADI is due to Wachspress and collaborators [39] [40] [32] [21]. A low-rank version of ADI, which is not considered here, was first proposed in 2000 by Penzl [35], and since then has become a very popular approach, see e.g. [15] [8] [7] [29]. For more details on the classical ADI method, we refer to the monograph [42].

The $j-$th iteration of the ADI method applied to equation (4.1) reads:

$$((K_1 + \omega_j M_1) \otimes M_2)\, s_{j-1/2} = r - (M_1 \otimes (K_2 - \omega_k M_2))\, s_{j-1}, \tag{4.4a}$$

$$(M_1 \otimes (K_2 + \gamma_j M_2))\, s_j = r - ((K_1 - \gamma_j M_1) \otimes M_2)\, s_{j-1/2}, \tag{4.4b}$$

where $\gamma_j, \omega_j \in \mathbb{R}$ are acceleration parameters, which will be discussed in the next section. Note that at each ADI iteration we need to solve two linear systems where the coefficient matrix has a Kronecker product structure, and this can be done efficiently by (2.11).

4.2.1. Convergence analysis. Let J denote the total number of ADI iterations, and let $e_J = s - s_J$, where s is the exact solution of (4.1), denote the final error. Then it holds

$$e_J = M^{-1/2} T_J M^{1/2} e_0, \tag{4.5}$$

where

$$T_J = \prod_{j=1}^{J} \left(\widetilde{K_1} - \gamma_j I\right)\left(\widetilde{K_1} + \omega_j I\right)^{-1} \otimes \left(\widetilde{K_2} + \gamma_j I\right)^{-1}\left(\widetilde{K_2} - \omega_j I\right), \tag{4.6}$$

with $M = M_1 \otimes M_2$, $\widetilde{K_1} = M_1^{-1/2} K_1 M_1^{-1/2}$ and $\widetilde{K_2} = M_2^{-1/2} K_2 M_2^{-1/2}$.

Given a vector $v \in \mathbb{R}^N$, its $M-$norm is defined as $\|v\|_M := \sqrt{v^T M v}$. Note that $\|v\|_M = \left\|M^{1/2} v\right\|$. With this definition, from (4.5) we infer that

$$\frac{\|e_J\|_M}{\|e_0\|_M} \leq \|T_J\|.$$

If $\Lambda\left(M_1^{-1} K_1\right) \subseteq [a, b]$ and $\Lambda\left(M_2^{-1} K_2\right) \subseteq [c, d]$, then it holds

$$\|T_J\| \leq \max_{\lambda \in [a,b],\, \mu \in [c,d]} \prod_{j=1}^{J} \left| \frac{\lambda - \gamma_j}{\lambda + \omega_j} \cdot \frac{\mu - \omega_j}{\mu + \gamma_j} \right|. \tag{4.7}$$

An explicit expression for the parameters γ_j, ω_j, $j = 1, \ldots, J$ which minimize the right-hand side of (4.7) is known [40]. When these parameters are selected the convergence behavior of ADI is well-understood. In particular, if we assume for simplicity that $[a, b] = [c, d]$, then the number of ADI iterations needed to ensure that $\|T_J\| \leq \epsilon$, for small enough ϵ, is

$$J = \left\lceil \frac{1}{\pi^2} \ln\left(4\frac{b}{a}\right) \ln\left(\frac{4}{\epsilon}\right) \right\rceil, \tag{4.8}$$

where $\lceil \cdot \rceil$ denotes the integer round toward positive infinity. We emphasize that the dependence of J on the spectral condition number of the matrices involved is logarithmic, and hence extremely mild. Moreover, once a tolerance ϵ is chosen, the number of ADI iterations can be selected a priori according to (4.8).

4.2.2. The algorithm. A simple trick to reduce the computational cost of each ADI iteration is to define

$$\tilde{s}_i = \begin{cases} (M_1 \otimes I_n)\, s_i & \text{for } i \in \mathbb{N}, \\ (I_n \otimes M_2)\, s_i & \text{for } i \notin \mathbb{N}. \end{cases}$$

Equations (4.4) now read

$$\begin{aligned} \left((K_1 + \omega_j M_1) \otimes I_n\right) \tilde{s}_{j-1/2} &= r - \left(I_n \otimes (K_2 - \omega_j M_2)\right) \tilde{s}_{j-1}, \\ \left(I_n \otimes (K_2 + \gamma_j M_2)\right) \tilde{s}_j &= r - \left((K_1 - \gamma_j M_1) \otimes I_n\right) \tilde{s}_{j-1/2}. \end{aligned}$$

We now summarize the steps of the ADI method.

Algorithm 2 ADI method

Fix the number of iterations J according to (4.8)
Compute the parameters ω_j, γ_j, $j = 1, \dots, J$.
Set $\tilde{s}_0 = 0$.
for $j = 1, \dots, J$ **do**
    Set $r_{j-1} = r - (I_n \otimes (K_2 - \omega_j M_2))\, \tilde{s}_{j-1}$.
    Solve $((K_1 + \omega_j M_1) \otimes I_n)\, \tilde{s}_{j-1/2} = r_{j-1}$.
    Set $r_{j-1/2} = r - ((K_1 - \gamma_j M_1) \otimes I_n)\, \tilde{s}_{j-1/2}$.
    Solve $(I_n \otimes (K_2 + \gamma_j M_2))\, \tilde{s}_j = r_{j-1/2}$.
end for
Set $s_J = (M_1 \otimes I_n)^{-1}\, \tilde{s}_J$.

4.2.3. ADI as a preconditioner. Our main interest is to apply ADI as a preconditioner. We observe that, if we take as initial guess $s_0 = 0$ then equality (4.5) can be rewritten as

$$s_J = M^{-1/2}\,(I - T_J)\,M^{1/2}\mathcal{P}^{-1}r =: \mathcal{P}_J^{-1}r.$$

Hence performing J iterations of the ADI method to the system $\mathcal{P}s = r$ is equivalent to multiply r by the matrix $\mathcal{P}_J^{-1} = M^{-1/2}\,(I - T_J)\,M^{1/2}\mathcal{P}^{-1}$. We emphasize that such equivalency is just a theoretical tool, and it is never used to actually compute s_J. We also emphasize that, since the number of iterations J is fixed by the chosen tolerance ϵ, the preconditioner $\mathcal{P}_J^{-1}$ does not change between different CG iterations.

In order to apply ADI as a preconditioner for CG, two issues has to be addressed. First, the CG method may break down if an arbitrary preconditioner is considered; in order to safely use the ADI method, we need to show that $\mathcal{P}_J$ is symmetric and positive definite. Second, the choice of the tolerance ϵ for ADI is crucial and has to be discussed. Indeed, a tolerance that is too strict yields unnecessary work, while a tolerance that is too loose may compromise the convergence of CG.

The following theorem addresses both issues. In particular, it presents a nice and simple upper bound for the spectral conditioning of $\mathcal{P}_J^{-1}\mathcal{A}$ in terms of ϵ and of the conditioning of the exactly preconditioned system $\mathcal{P}^{-1}\mathcal{A}$. A proof of this theorem, which generalizes the results of [40, Section 3], can be found in [42]. We give a proof of this theorem in our notation, to keep the present manuscript as self-contained as possible.

THEOREM 1.

The ADI preconditioner $\mathcal{P}_J$ (with optimal parameters) is positive definite. Moreover, if $\|T_J\| \leq \epsilon$, then it holds

$$\kappa\left(\mathcal{P}_J^{-1}\mathcal{A}\right) \leq \left(\frac{1+\epsilon}{1-\epsilon}\right)\kappa\left(\mathcal{P}^{-1}\mathcal{A}\right). \tag{4.9}$$

Proof.

We note that $\mathcal{P}_J$ is symmetric and positive definite if and only if the same holds true for

$$M^{1/2}\mathcal{P}_J^{-1}M^{1/2} = (I - T_J)\,\widetilde{\mathcal{P}}^{-1}, \tag{4.10}$$

with $\widetilde{\mathcal{P}} = M^{-1/2}\mathcal{P}M^{-1/2}$. We have already observed that T_J is symmetric. Moreover, $I - T_J$ is positive definite, since $\|T_J\| < 1$. The last inequality follows from (4.7); indeed

$p=1$	$p=2$	$p=3$	$p=4$	$p=5$	$p=6$	$p=7$	$p=8$
$3\cdot 10^5$	$3\cdot 10^5$	$4\cdot 10^5$	$4\cdot 10^5$	$1\cdot 10^6$	$2\cdot 10^6$	$2\cdot 10^6$	$3\cdot 10^6$

TABLE 4.1
Spectral condition number of $M_1^{-1}K_1$, with $h=2^{-5}$.

if we take any set of parameters that satisfy $0<\gamma_j,\omega_j\leq \min\{a,c\}$, $j=1,\dots,J$, then each factor of the product is strictly smaller than 1. Since the optimal parameters minimize $\|T_J\|$, also in this case we have $\|T_J\|<1$.

We observe that $\widetilde{\mathcal{P}}=\widetilde{K_1}\otimes I+I\otimes\widetilde{K_2}$ and $I-T_J$ share the same set of eigenvectors, and since they are both symmetric and positive definite, the product in (4.10) is again symmetric and positive definite.

We now turn on inequality (4.9). Matrix $\mathcal{P}_J^{-1}\mathcal{A}$ is similar to

$$\widetilde{\mathcal{P}}^{1/2}M^{1/2}\mathcal{P}_J^{-1}\mathcal{A}M^{-1/2}\widetilde{\mathcal{P}}^{-1/2}=(I-T_J)\left(\widetilde{\mathcal{P}}^{-1/2}M^{-1/2}\mathcal{A}M^{-1/2}\widetilde{\mathcal{P}}^{-1/2}\right),$$

where we used the fact that $\widetilde{\mathcal{P}}^{1/2}$ and $I-T_J$ commute, since they share the same set of eigenvectors.

It holds

$$\begin{aligned}\lambda_{\max}\left(\mathcal{P}_J^{-1}\mathcal{A}\right)&=\lambda_{\max}\left((I-T_J)\left(\widetilde{\mathcal{P}}^{-1/2}M^{-1/2}\mathcal{A}M^{-1/2}\widetilde{\mathcal{P}}^{-1/2}\right)\right)\\&\leq(1+\epsilon)\,\lambda_{\max}\left(\widetilde{\mathcal{P}}^{-1/2}M^{-1/2}\mathcal{A}M^{-1/2}\widetilde{\mathcal{P}}^{-1/2}\right)\\&=(1+\epsilon)\,\lambda_{\max}\left(\mathcal{P}^{-1}\mathcal{A}\right).\end{aligned}\tag{4.11}$$

With analogous computations one can show that

$$\lambda_{\min}\left(\mathcal{P}_J^{-1}\mathcal{A}\right)\geq(1-\epsilon)\,\lambda_{\min}\left(\mathcal{P}^{-1}\mathcal{A}\right).\tag{4.12}$$

Combining (4.12) and (4.11), inequality (4.9) is proved. □

4.2.4. Computational cost. Lines 1 and 2 of Algorithm 2, namely the computation of J and of the optimal parameters, require an estimate of the minimum and maximum eigenvalue of $M_1^{-1}K_1$ and $M_2^{-1}K_2$. Since all the matrices involved are symmetric, positive definite and banded, this task can be achieved inexpensively, e.g. with a few iterations of the power method.

The computational effort of one full ADI iteration consists in $2n$ matrix-vector products (lines 5 and 7 of Algorithm 2) and the solution of $2n$ linear systems (lines 6 and 8), both of which involve banded matrices of order n. Since lines 5 and 7 also require a vector update of order n^2, the total cost of a single ADI iteration is roughly $(16p+10)\,n^2$ FLOPs.

The cost of a single iteration has to be multiplied for the number iterations J, given by (4.8), in order to obtain the total cost of ADI. For IGA mass and stiffness matrices, it holds $\dfrac{b}{a}=\kappa\left(M_1^{-1}K_1\right)\leq c_p h^{-2}\approx c_p n^2$, where c_p is a constant that depends only on p. From inverse estimates for polynomials it follows $c_p\leq p^4$, however no accurate estimate of c_p is known yet. Hence, we seek numerical evidence of the dependence of $\kappa\left(M_1^{-1}K_1\right)$ w.r.t. p. The results, obtained for $h=2^{-5}$, are reported in Table 4.1. We can see that the growth of the conditioning is not dramatic, with a growth which is weaker than the bound above and, since we are only interested in its logarithm, we can conclude that the number of ADI iterations is robust w.r.t. p.

In conclusion the total cost of ADI is roughly

$$\left[\frac{1}{\pi^2}\ln\left(4c_p n^2\right)\ln\left(\frac{4}{\epsilon}\right)+1\right](16p+10)\,n^2$$

FLOPs. We observe that the cost of the preconditioner is (almost) linear w.r.t. to the number of degrees-of-freedom, and very robust w.r.t. p. Indeed, this cost has a milder dependence on p than the cost of a matrix-vector product with $\mathcal{A}$, which is $2(2p+1)^2n^2$ FLOPs.

5. The 3D case. When $d=3$, equation (3.2) takes the form

$$\left(K_1\otimes M_2\otimes M_3+M_1\otimes K_2\otimes M_3+M_1\otimes M_2\otimes K_3\right)s=r. \tag{5.1}$$

We consider generalizations of the approaches detailed in the previous section, namely the FD direct method and the ADI iterative method.

Other approaches, which however rely on a low-rank approximation of the right-hand side, can be found in [25], [30] and [1].

5.1. The 3D fast diagonalization method. The direct method presented in section 4.1 admits a straightforward generalization to the 3D case (see also [31], where the Bartels-Stewart approach for the nonsymmetric case is extended to 3D problems). We consider the generalized eigendecompositions

$$K_1U_1=M_1U_1D_1,\qquad K_2U_2=M_2U_2D_2,\qquad K_3U_3=M_3U_3D_3, \tag{5.2}$$

with D_1, D_2, D_3 diagonal matrices and

$$U_1^TM_1U_1=I,\qquad U_2^TM_2U_2=I,\qquad U_3^TM_3U_3=I.$$

Then, (5.1) can be factorized as

$$\left(U_1\otimes U_2\otimes U_3\right)^{-1}\left(D_1\otimes I\otimes I+I\otimes D_2\otimes I+I\otimes I\otimes D_3\right)\left(U_1\otimes U_2\otimes U_3\right)^{-T}s=r,$$

which suggests the following algorithm.

Algorithm 3 FD direct method (3D)

1: Compute the generalized eigendecompositions (5.2)
2: Compute $\tilde{r}=(U_1\otimes U_2\otimes U_3)r$
3: Compute $\tilde{s}=\left(D_1\otimes I\otimes I+I\otimes D_2\otimes I+I\otimes I\otimes D_3\right)^{-1}\tilde{r}$
4: Compute $s=(U_1\otimes U_2\otimes U_3)^T\tilde{s}$

5.1.1. Computational cost. Lines 1 and 3 require $O(n^3)$ FLOPs. Lines 2 and 4, as can be seen by nested applications of formula (2.13), are equivalent to performing a total of 6 products between dense matrices of size $n\times n$ and $n\times n^2$. Thus, neglecting lower order terms the overall computational cost of Algorithm 3 is $12n^4$ FLOPs.

The direct method is even more appealing in the 3D case than it was in the 2D case, for at least two reasons. First, the computational cost associated with the preconditioner setup, that is the eigendecomposition, is negligible. This means that the main computational effort of the method consists in a few (dense) matrix-matrix products, which are level 3 BLAS operations and typically yield high efficiency thanks

to a dedicated implementation on modern computers by optimized usage of the memory cache hierarchy [24, Chapter 1]. Second, in a preconditioned CG iteration the cost for applying the preconditioner has to be compared with the cost of the residual computation (a matrix-vector product with $\mathcal{A}$) which can be quantified in approximately $2(2p+1)^3n^3$ for 3D problems, resulting in a FLOPs ratio of the preconditioner application to residual computation of $(3n)/(4p^3) \approx n/p^3$. For example, if $N = 256^3$ and $p = 4$, the preconditioner requires only 3 times more FLOPs than the residual computation, while for degree $p = 6$ the matrix-vector product is even more costly than the preconditioner itself. However in numerical tests we will see that, for all cases of practical interest in 3D, the computational time used by the preconditioner application is far lower that the residual computation itself. This is because the computational time depends not only on the FLOPs count but also on the memory usage and, as mentioned above, dense matrix-matrix multiplications greatly benefit of modern computer architecture. This approach will show largely higher performance than the alternative ADI approach we have considered.

5.2. Three-variable ADI. Despite the clear advantages presented by the direct method discussed in the previous section, for the sake of comparison we also consider ADI. Indeed ADI may benefit from a lower FLOPs counting than the direct solver, for large n and low p, and being an iterative solver we can optimize the target precision as needed by the preconditioning step.

However, the ADI extension to the 3D case is not straightforward. Here we follow the iterative scheme proposed by Douglas in [20] (see also [41] for a different approach) to solve (5.1):

$$(\mathcal{K}_1 + \omega_j M)\, s_{j-2/3} = 2r - (\mathcal{K}_1 + 2\mathcal{K}_2 + 2\mathcal{K}_3 - \omega_j M)\, s_{j-1}, \tag{5.3a}$$

$$(\mathcal{K}_2 + \omega_j M)\, s_{j-1/3} = \mathcal{K}_2 s_{j-1} + \omega_j M s_{j-2/3}, \tag{5.3b}$$

$$(\mathcal{K}_3 + \omega_j M)\, s_j = \mathcal{K}_3 s_{j-1} + \omega_j M s_{j-1/3}, \tag{5.3c}$$

where $\mathcal{K}_1 = K_1 \otimes M_2 \otimes M_3$, $\mathcal{K}_2 = M_1 \otimes K_2 \otimes M_3$, $\mathcal{K}_3 = M_1 \otimes M_2 \otimes K_3$, $M = M_1 \otimes M_2 \otimes M_3$, and the ω_j are real positive parameters. After J steps, reasoning as in the 2D case, we can derive an expression for the error e_J similar to (4.5):

$$e_J = M^{-1/2} T_J M^{1/2} e_0,$$

where T_J is a symmetric positive definite matrix that depends on $\omega_1, \dots, \omega_J$. Hence, the relative error in the $M-$norm is bounded by the euclidean norm of T_J.

We assume for simplicity that $\Lambda\left(M_1^{-1}K_1\right) = \Lambda\left(M_2^{-1}K_2\right) = \Lambda\left(M_3^{-1}K_3\right) = [a, b]$. Then it can be shown that

$$\|T_J\| \leq \max_{\lambda_1,\lambda_2,\lambda_3 \in [a,b]} \left| \prod_{j=1}^{J} \left(1 - 2\omega_j^2 \frac{\lambda_1 + \lambda_2 + \lambda_3}{(\omega_j + \lambda_1)\,(\omega_j + \lambda_2)\,(\omega_j + \lambda_3)} \right) \right| =: \rho_J(\omega_1, \dots, \omega_J).$$

Clearly we are interested in choosing the parameters ω_j, $j = 1, \dots, J$ so that the right-hand side of the above inequality is minimized. However, unlike in the 2D case, no expression for the solution of such minmax problem is known, and hence we cannot rely on an optimal choice for the parameters. This makes the ADI approach less appealing than in the 2D case. The suboptimal choice proposed in [20] still guarantees that the number of iterations J_0 needed to ensure that $\|T_{J_0}\| \leq \epsilon$ satisfies

$$J_0 \;\approx\; 1.16\, \ln\left(\frac{b}{a}\right) \ln\left(\epsilon^{-1}\right) = O\left(\ln\left(\frac{b}{a}\right) \ln\left(\epsilon^{-1}\right) \right), \tag{5.4}$$

Note that also in the 2D case we have that the number of iterations required to achieve convergence is $O\left(\ln\left(\frac{b}{a}\right)\ln\left(\epsilon^{-1}\right)\right)$ (cf. (4.8)). However, in 3D the constant hidden in this asymptotic estimate is significantly greater than in the 2D case.

In fact, our numerical experience indicates that typically the condition $\|T_J\| \leq \epsilon$ is satisfied after much fewer iterations than (5.4) would suggest. To avoid unnecessary iterations, we introduce a different stopping criterion, based on the evaluation of the right hand side of (5.3). More precisely, we compute the parameters $\omega_1, \ldots, \omega_{J_0}$ according to [20] but then perform only the first $J \leq J_0$ iterations, where J is the smallest index such that $\rho_J(\omega_1, \ldots, \omega_J) \leq \epsilon$.

Another possible approach is to discard the suboptimal choice of [20], and select the parameters $\omega_1, \ldots, \omega_J$ following a greedy strategy, in the spirit of the approach proposed in [35] for the two-variables ADI method. After iteration k, we compute

$$(\lambda_1^*, \lambda_2^*, \lambda_3^*) = \underset{\lambda_1,\lambda_2,\lambda_3 \in [a,b]}{\operatorname{argmax}} \left| \prod_{j=1}^{k} \left(1 - 2\omega_j^2 \frac{\lambda_1 + \lambda_2 + \lambda_3}{(\omega_j + \lambda_1)(\omega_j + \lambda_2)(\omega_j + \lambda_3)} \right) \right|, \tag{5.5}$$

and then ω_{k+1} is defined as the nonnegative number that minimizes the error at $(\lambda_1^*, \lambda_2^*, \lambda_3^*)$, that is

$$\omega_{k+1} = \underset{\omega \geq 0}{\operatorname{argmin}} \left| 1 - 2\omega^2 \frac{\lambda_1^* + \lambda_2^* + \lambda_3^*}{(\omega + \lambda_1^*)(\omega + \lambda_2^*)(\omega + \lambda_3^*)} \right|. \tag{5.6}$$

To reduce the computational cost of one ADI iteration, we multiply the first equation of (5.3) by $(I_n \otimes M_2 \otimes M_3)^{-1}$, the second by $(M_1 \otimes I_n \otimes M_3)^{-1}$ and the third by $(M_1 \otimes M_2 \otimes I_n)^{-1}$. After some algebraic manipulation we obtain

$$\begin{aligned}
((K_1 + \omega_j M_1) \otimes I_n \otimes I_n)\, s_j^* &= \tilde{r} - ((K_1 - \omega_j M_1) \otimes I_n \otimes I_n)\, s_{j-1} - 2\,(M_1 \otimes I_n \otimes I_n) \\
&\quad \cdot \left(I_n \otimes M_2^{-1} K_2 \otimes I_n + I_n \otimes I_n \otimes M_3^{-1} K_3\right) s_{j-1}, \qquad (5.7a) \\
(I_n \otimes (K_2 + \omega_j M_2) \otimes I_n)\, s_j^{**} &= (I_n \otimes M_2 \otimes I_n) \\
&\quad \cdot \left(\left(I_n \otimes M_2^{-1} K_2 \otimes I_n\right) s_{j-1} + \omega_j s_j^*\right), \qquad (5.7b) \\
(I_n \otimes I_n \otimes (K_3 + \omega_j M_3))\, s_j &= (I_n \otimes I_n \otimes M_3) \\
&\quad \cdot \left(\left(I_n \otimes I_n \otimes M_3^{-1} K_3\right) s_{j-1} + \omega_j s_j^{**}\right), \qquad (5.7c)
\end{aligned}$$

where $\tilde{r} = 2\,(I_n \otimes M_2 \otimes M_3)^{-1}\, r$. Note that the vectors

$$u_j := \left(I_n \otimes M_2^{-1} K_2 \otimes I_n\right) s_{j-1}, v_j := \left(I_n \otimes I_n \otimes M_3^{-1} K_3\right) s_{j-1},$$

which both appear twice in (5.7), need to be computed only once. We consider one last trick to save some computational cost. Let

$$b_j := \left(I_n \otimes I_n \otimes M_3^{-1} K_3\right) s_{j-1} + \omega_j s_j^{**},$$

then

$$v_{j+1} = (I_n \otimes I_n \otimes M_3)^{-1}\, (I_n \otimes I_n \otimes (K_3 + \omega_j M_3))\, s_j - \omega_j s_j = b_j - \omega_j s_j.$$

where the last equality is a consequence of equation (5.7c). Hence we can use the known vectors b_j and s_j to inexpensively compute v_{j+1}. We summarize all these considerations in Algorithm 4.

As for the previous methods, we are mainly interested in using ADI as a preconditioner. It can be shown that Theorem 1 holds also in the 3D case. We do not report the details, as the arguments used in the proof are the same as in the 2D case. This means that the 3D ADI method can be used as a preconditioner for CG and that relation (4.9) guides in choosing the inner tolerance ϵ.

Algorithm 4 3D ADI method

1: Compute all the eigenvalues of $M_1^{-1}K_1$, $M_2^{-1}K_2$ and $M_3^{-1}K_3$.
2: Compute the number of iterations J and parameters $\omega_1, \dots, \omega_J$ such that $\rho_J(\omega_1, \dots, \omega_J) \leq \epsilon$.
3: Compute $\tilde{r} = 2\left(I_n \otimes M_2 \otimes M_3\right)^{-1} r$.
4: Set $s_0, v_0 = 0$.
5: **for** $j = 1, \dots, J$ **do**
6: Solve $\left(I_n \otimes M_2 \otimes I_n\right) u_j = \left(I_n \otimes K_2 \otimes I_n\right) s_{j-1}$
7: Compute the right-hand side of (5.7a):
8: $r_j^* = \tilde{r} - \left((K_1 - \omega_j M_1) \otimes I_n \otimes I_n\right) s_{j-1} - \left(2M_1 \otimes I_n \otimes I_n\right)\left(u_j + v_j\right)$
9: Solve $\left((K_1 + \omega_j M_1) \otimes I_n \otimes I_n\right) s_j^* = r_j^*$
10: Compute the right-hand side of (5.7b): $r_j^{**} = \left(I_n \otimes M_2 \otimes I_n\right)\left(u_j + \omega_j s_j^*\right)$
11: Solve $\left(I_n \otimes (K_2 + \omega_j M_2) \otimes I_n\right) s_j^{**} = r_j^{**}$
12: Set $b_j = v_j + \omega_j s_j^{**}$
13: Compute the right-hand side of (5.7c): $r_j = \left(I_n \otimes I_n \otimes M_3\right) b_j$
14: Solve $\left(I_n \otimes I_n \otimes (K_3 + \omega_j M_3)\right) s_j = r_j$
15: Set $v_{j+1} = b_j - \omega_j s_j$
16: **end for**

5.2.1. Computational cost. At each iteration, the main computational effort is represented by the solution of four linear systems (lines 5, 7, 9 and 12) and five matrix products (lines 5, 6, 8, and 11). As always, by exploiting the Kronecker structure of the matrices involved, each of these computations can be performed at a cost of $2(2p+1)n^3$ FLOPs. A careful analysis reveals that the total cost of a single ADI iteration is $(36p + 29)\, n^3$ FLOPs, where we neglected terms of order lower than n^3.

Unlike in the 2D case, the number of iterations J is not known a priori. However, if we consider J_0 in (5.4) as an upper bound for the number of iterations, we can bound the total computational cost can be bounded roughly by

$$1.16\ \ln\left(c_p n^2\right) \ln\left(\epsilon^{-1}\right) (36p + 29)\, n^3. \tag{5.8}$$

where, as in the 2D case, we replaced $\dfrac{b}{a}$ with $c_p n^2$. We observe that asymptotically the complexity of ADI is almost $\mathrm{O}(n^3)$, that is almost linear w.r.t. the number of degrees-of-freedom. On the other hand, a closer look at (5.8) reveals that the number of FLOPs required may be actually quite large, even for small or moderate p. We remark that the sequential nature of the ADI iteration makes it less suited for parallelization than the FD method.

6. Application to multi-patch problems. To enhance flexibility in geometry representation, typically multi-patch parametrizations are adopted in isogeometric analysis. This means that the domain of interest Ω is the union of patches Ω_i such that $\Omega_i = \mathbf{F}_i([0,1]^d)$, and each $\mathbf{F}_i$ is a spline (or NURBS) parametrization. Typically $\Omega_i \cap \Omega_j$ is an empty set, or a vertex, or the full common edge or the full common face (when $d = 3$) of the patches. Furthermore we assume that the meshes are conforming, that is for each patch interface $\Omega_i \cap \Omega_j$ the isogeometric functions on Ω_i and the ones on Ω_j generate the same trace space. See, e.g., [4] for more details.

For such a configuration, we can easily combine the approaches discussed in the previous sections with an overlapping Schwarz preconditioner. For that, we need to

further split Ω into overlapping subdomains. We choose the subdomains as pairs of neighboring patches merged together. Precisely, let N_s denotes the total number of interfaces between neighboring patches. We define

$$\Theta_i = \Omega_{i_1} \cup \Omega_{i_2}, \qquad i = 1, \dots, N_s,$$

where Ω_{i_1} and Ω_{i_2} are the patches which share the $i-$th interface. It holds $\Omega = \bigcup_i \Theta_i$.

Now let R_i be the rectangular restriction matrix on the degrees-of-freedom associated with the $i-$th subdomain, and let $\mathcal{A}_i = R_i \mathcal{A} R_i^T$.The (exact) additive Schwarz preconditioner is

$$\mathcal{P}_{EAS} = \sum_{i=1}^{N_s} R_i^T \mathcal{A}_i^{-1} R_i, \tag{6.1}$$

and its inexact variant

$$\mathcal{P}_{IAS} = \sum_{i=1}^{N_s} R_i^T \widetilde{\mathcal{A}}_i^{-1} R_i, \tag{6.2}$$

where each $\widetilde{A}_i^{-1}$ is a suitable approximation of $\mathcal{A}_i^{-1}$. Each $\mathcal{A}_i$ represents the system matrix of a discretized Poisson problem on Θ_i. A crucial observation is that, under the conforming assumption, Θ_i can be considered a single-patch domain. Thus, it is possible to construct a preconditioner of the form (3.1) for $\mathcal{A}_i$, which we denote with $\mathcal{P}_i$. Then $\mathcal{P}_i$ can be used to construct $\widetilde{\mathcal{A}}_i^{-1}$ (we can have $\widetilde{\mathcal{A}}_i = \mathcal{P}_i$, or $\widetilde{\mathcal{A}}_i^{-1}$ may represent a fixed number of iteration of some iterative method preconditioned by $\mathcal{P}_i$). The proposed approach is somewhat unusual in the context of domain decomposition methods. Indeed, it is more common to split the domain Ω into a large number of small subdomains, so that local problems can be efficiently solved by parallel architectures. Here, on the other hand, the subdomains are chosen so that the basis functions of the local problems have a tensor structure that can be exploited by our preconditioner. The efficiency of such preconditioners, which is demonstrated numerically in the next section, makes it feasible to work with local problems whose size is comparable with that of the whole system. Finally, we remark that the large overlap between neighboring subdomains ensures that the outer iteration converges fast, independently of p and h.

Extension of this approach to nonconforming discretizations would require the use of nonconforming DD preconditioners (e.g., [28]) instead of an overlapping Schwarz preconditioner.

7. Numerical experiments. We now numerically show the potential of the approaches described in Sections 4 and 5. All the algorithms are implemented in MATLAB Version 8.5.0.197613 (R2015a), with the toolbox GEOPDES [16], on a Intel Xeon i7-5820K processor, running at 3.30 GHz, and with 64 GB of RAM. Although the Sylvester-based approaches are very suited for parallelization, particularly the FD method, here we benchmark sequential execution and use only one core for the experiments.

We give a few technical details on how the methods were implemented. For the FD method, we used the MATLAB function `eig` to compute the generalized eigendecomposition (4.3) and (5.2). For 2D ADI, the number of iterations was set according

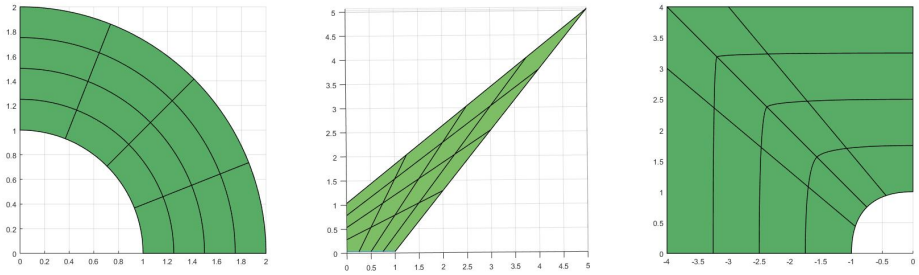

FIG. 7.1. *Quarter of ring, stretched square and plate with hole domains*

to (4.8). The extreme eigenvalues of $M_1^{-1}K_1$ and $M_2^{-1}K_2$, which are required for computing the optimal parameters derived in [40], were approximated using 10 iterations of the (direct and inverse) power method. For 3D ADI, the eigenvalue computation necessary to select the parameters was again performed using `eig`. In both 2D and 3D, at each ADI iteration the linear systems were solved using MATLAB's direct solver "backslash". Finally, in 3D algorithms the products involving Kronecker matrices were performed using the function from the free MATLAB toolbox Tensorlab [37].

Although in many of the problems considered here the matrix pencils (K_1, M_1), (K_2, M_2) (and in the 3D case (K_3, M_3)) coincide, we never exploit this fact in our tests. For example, in line 1 of Algorithm 1 we always compute two eigendecompositions even if in the current problem we have $M_1 = M_2$ and $K_1 = K_2$. In this way, the computational effort reflects the more general case in which such matrices are different.

7.1. 2D experiments. We start by considering 2D problems, and observe the performance of the FD and ADI methods on four test problems with different geometries: a square, a quarter of ring, a stretched square and plate with hole. In all problems we set K as the identity matrix, since according to Proposition 1 the presence of coefficients and of a nontrivial geometry have an analogous impact on the difficulty of the problem.

The square domain is simply $[0,1]^2$, and the other domains are shown in Figure 7.1. In the case of the plate with hole we chose the same parametrization considered in [13, Section 4.2]. In particular, two control points are placed in the same spacial location, namely the left upper corner, and this creates a singularity in the Jacobian of $\boldsymbol{F}$. Thus, in this case the bound provided by Proposition 1 becomes $\kappa\left(\mathcal{P}^{-1}\mathcal{A}\right) \leq +\infty$, and it is hence useless. In principle, our approaches could perform arbitrarily bad and this problem is indeed intended to test their performance in this unfavorable case. In all problems, except the last one, the system $\mathcal{A}u = b$ represents the discretization of problem (2.1), with $f = 2\left(x^2 - x\right) + 2\left(y^2 - y\right)$. For the plate with hole domain, we considered $f = 0$ and mixed boundary conditions.

We start by considering the problem on the square. As already said, in this case $\mathcal{P} = \mathcal{A}$ and we can directly apply the considered method to the system $\mathcal{A}u = b$. This is not a realistic case but serves as a preliminary check on the proposed theory and implementations. Results are shown in Table 7.1; in the upper part we report the CPU time for the FD method, while in the lower part we report the number of iterations and CPU time for ADI, whose tolerance was set to 10^{-8}, for different values of h and p.

We observe that the computation time for the direct method is substantially

	FD Direct Solver Time (sec)					
h^{-1}	$p=1$	$p=2$	$p=3$	$p=4$	$p=5$	$p=6$
512	0.18	0.16	0.16	0.16	0.16	0.15
1024	1.52	1.17	1.02	0.94	0.95	0.89
2048	10.62	10.04	11.70	9.21	7.64	6.68
4096	72.73	71.72	127.42	108.91	68.91	83.83
8192	511.33	511.27	1145.04	1030.96	515.40	856.62

	ADI Iterations / Time (sec)					
h^{-1}	$p=1$	$p=2$	$p=3$	$p=4$	$p=5$	$p=6$
512	29 / 0.34	28 / 0.33	29 / 0.37	30 / 0.40	31 / 0.43	32 / 0.45
1024	31 / 1.72	31 / 1.56	32 / 1.64	33 / 1.82	34 / 1.96	35 / 2.05
2048	34 / 8.39	34 / 11.61	35 / 8.39	36 / 10.42	37 / 10.62	37 / 9.23
4096	37 / 37.25	37 / 52.59	37 / 37.48	38 / 40.42	39 / 43.03	40 / 40.25
8192	40 / 160.91	39 / 218.11	40 / 161.57	41 / 173.67	42 / 186.61	43 / 172.50

TABLE 7.1
Square domain. Performance of FD (upper table) and of ADI (lower table)

h^{-1}	$p=1$	$p=2$	$p=3$	$p=4$	$p=5$	$p=6$
512	$3.0 \cdot 10^{-9}$	$7.0 \cdot 10^{-9}$	$6.8 \cdot 10^{-9}$	$6.4 \cdot 10^{-9}$	$6.7 \cdot 10^{-9}$	$6.7 \cdot 10^{-9}$
1024	$7.7 \cdot 10^{-9}$	$7.2 \cdot 10^{-9}$	$5.4 \cdot 10^{-9}$	$6.2 \cdot 10^{-9}$	$6.3 \cdot 10^{-9}$	$5.8 \cdot 10^{-9}$
2048	$7.5 \cdot 10^{-9}$	$6.0 \cdot 10^{-9}$	$4.9 \cdot 10^{-9}$	$6.0 \cdot 10^{-9}$	$5.1 \cdot 10^{-9}$	$9.7 \cdot 10^{-9}$
4096	$6.9 \cdot 10^{-9}$	$4.9 \cdot 10^{-9}$	$7.6 \cdot 10^{-9}$	$9.2 \cdot 10^{-9}$	$7.8 \cdot 10^{-9}$	$8.4 \cdot 10^{-9}$
8192	$6.1 \cdot 10^{-9}$	$8.2 \cdot 10^{-9}$	$7.5 \cdot 10^{-9}$	$9.3 \cdot 10^{-9}$	$7.9 \cdot 10^{-9}$	$8.2 \cdot 10^{-9}$

TABLE 7.2
ADI relative error $\|e_J\|_M / \|e_0\|_M$ at the final iteration for problems in Table 7.1

independent w.r.t. p; fluctuations in time appearing in the finer discretization levels are due to the `eig` function, which constitute the main computational effort of the method in our implementation. Similarly, computation times for ADI do not change significantly by varying p, and fluctuations are due to MATLAB's direct solver.

Concerning the dependence on h, based on the analysis of computational cost we expect ADI to perform better than FD for small enough h. That is indeed what we can see in the experimental results; however, ADI starts outperforming FD for a very small value of h, corresponding roughly to 16 million degrees-of-freedom. Moreover, we emphasize that the CPU times of the two methods are comparable for all the discretization levels considered.

The number of iterations of ADI is determined a priori, according to (4.8), and no a posteriori stopping criterion is considered. Table 7.2 reports the relative error $\|e_J\|_M / \|e_0\|_M$ for all cases considered in Table 7.1: observe that in all cases, this value is below the prescribed tolerance 10^{-8} and at the same time, never smaller than $3 \cdot 10^{-9}$, showing that (4.8) is indeed a good choice.

We now turn to the first two problems with nontrivial geometry, namely the quarter of ring and the stretched square, and employ FD and ADI as preconditioners for CG (represented respectively by matrices $\mathcal{P}$ and $\mathcal{P}_J$). For both problems, we set $\epsilon = 10^{-1}$ as tolerance for the ADI preconditioner. We remark that a slightly better performance could be obtained by an adaptive choice of ϵ, as described in [42, Chapter 3]. However, we did not implement this strategy.

To better judge the efficiency of the Sylvester-based approaches, we compare the results with those obtained when using a preconditioner based on the Incomplete Cholesky (IC) factorization (implemented by the MATLAB function `ichol`). To

	CG + $\mathcal{P}$ Iterations / Time (sec)			
h^{-1}	$p=2$	$p=3$	$p=4$	$p=5$
128	25 / 0.04	25 / 0.06	25 / 0.07	25 / 0.09
256	25 / 0.20	25 / 0.26	25 / 0.34	25 / 0.40
512	26 / 1.13	26 / 1.36	26 / 1.62	26 / 2.00
1024	26 / 7.30	26 / 8.13	26 / 9.09	26 / 10.52

	CG + $\mathcal{P}_J$ Iterations / Time (sec)			
h^{-1}	$p=2$	$p=3$	$p=4$	$p=5$
128	25 (5) / 0.12	25 (5) / 0.14	25 (5) / 0.17	25 (5) / 0.19
256	26 (5) / 0.49	26 (5) / 0.54	26 (6) / 0.76	25 (6) / 0.82
512	27 (6) / 2.39	27 (6) / 2.68	26 (6) / 2.94	26 (6) / 3.38
1024	27 (6) / 9.94	27 (6) / 11.00	27 (7) / 14.15	27 (7) / 16.15

	CG + IC Iterations / Time (sec)			
h^{-1}	$p=2$	$p=3$	$p=4$	$p=5$
128	65 / 0.12	49 / 0.18	40 / 0.21	33 / 0.28
256	130 / 0.92	98 / 1.16	80 / 1.51	65 / 1.87
512	264 / 7.94	198 / 9.47	160 / 11.42	128 / 13.29
1024	533 / 64.54	399 / 75.22	324 / 89.29	262 / 103.26

TABLE 7.3

Quarter of ring domain. Performance of CG preconditioned by FD (upper table), by ADI (middle table) and by Incomplete Cholesky (lower table).

improve the performance of this approach, we considered some preliminary reorderings of $\mathcal{A}$, namely those implemented by the MATLAB functions `symrcm`, `symamd` and `colperm`. The reported results refer to the `symrcm` reordering, which yields the best performance. We remark that incomplete factorizations have been considered as preconditioners for IGA problems in [11], where the authors observed that this approach is quite robust w.r.t. p.

In Tables 7.3 and 7.4, we report the total computation time and the number of CG iterations for both problems; when ADI is used, we also report the number of iterations performed at each application of the preconditioner. Here and throughout, the computation time includes the time needed to setup the preconditioner. Note that the considered values of the mesh size are larger than in the square case. Indeed, while in the square case we need to store only the blocks M_1, M_2, K_1 and K_2, in the presence of nontrivial geometry the whole matrix $\mathcal{A}$ has to be stored, and this is unfeasible for our computer resources when h is too small. Below we report some comments on the numerical results for the two geometries: the quarter of ring and the stretched square.

- For both the ADI and the FD preconditioners, the number of CG iterations is practically independent on p and slightly increases as the mesh is refined, but stays uniformly bounded according to Proposition 1. Moreover, in both approaches the computation times depend on p only mildly.
- The inexact application of $\mathcal{P}$ via the ADI method does not significantly affect the number of CG iterations. Moreover, the number of inner ADI iterations is roughly the same in all considered cases.
- The overall performance obtained with FD is slightly better than with ADI for all the considered discretization levels; if finer meshes are considered, ADI should eventually outperform FD.
- Interestingly, in the IC approach the number of CG iterations decreases for

	CG + $\mathcal{P}$ Iterations / Time (sec)				
h^{-1}	$p=1$	$p=2$	$p=3$	$p=4$	$p=5$
128	58 / 0.08	61 / 0.08	62 / 0.12	61 / 0.16	61 / 0.21
256	64 / 0.39	66 / 0.48	66 / 0.61	66 / 0.82	66 / 0.98
512	69 / 2.58	70 / 2.83	69 / 3.33	69 / 4.01	69 / 4.88
1024	73 / 18.45	73 / 18.56	72 / 20.88	72 / 23.46	71 / 26.94

	CG + $\mathcal{P}_J$ Iterations / Time (sec)				
h^{-1}	$p=1$	$p=2$	$p=3$	$p=4$	$p=5$
128	58 (5) / 0.25	61 (5) / 0.28	61 (5) / 0.33	61 (5) / 0.40	61 (5) / 0.48
256	65 (5) / 0.94	65 (5) / 1.23	65 (5) / 1.35	65 (6) / 1.87	66 (6) / 2.14
512	69 (6) / 5.76	70 (6) / 6.27	69 (6) / 6.85	69 (6) / 7.71	69 (6) / 8.85
1024	74 (6) / 28.53	73 (6) / 27.17	73 (6) / 29.25	72 (7) / 37.47	72 (7) / 42.71

	CG + IC Iterations / Time (sec)				
h^{-1}	$p=1$	$p=2$	$p=3$	$p=4$	$p=5$
128	50 / 0.07	38 / 0.07	26 / 0.11	22 / 0.14	20 / 0.21
256	102 / 0.46	88 / 0.64	53 / 0.68	44 / 0.94	38 / 1.24
512	213 / 3.91	232 / 7.17	115 / 5.73	89 / 6.99	76 / 8.45
1024	439 / 33.68	780 / 93.07	248 / 47.56	181 / 50.80	153 / 63.52

TABLE 7.4

Stretched square domain. Performance of CG preconditioned by FD (upper table), by ADI (middle table) and by Incomplete Cholesky (lower table).

higher p. On the other hand, the CPU time still increases due to the greater computational cost of forming and applying the preconditioner.

- For small enough h, both the ADI and the FD preconditioners yield better performance, in terms of CPU time, than the IC preconditioner.

Finally, in Table 7.5 we report the results for the plate with hole domain. As expected, the performance of the Sylvester-based preconditioners in this case is much worse than in the previous cases, and in particular they are not robust neither w.r.t. h nor w.r.t. p. One can introduce modifications of $\mathcal{P}$ that significantly improve the conditioning of the preconditioned system, however we postpone this investigation to a further work. Interestingly, however, if we compare the results with the ones obtained with the IC preconditioner, we see that computation times are still comparable with those relative to the FD preconditioner for all discretization levels. In conclusion, even in most penalizing case among those considered, the proposed preconditioner is competitive with a standard one.

7.2. 3D experiments. As in the 2D case, we first consider a domain with trivial geometry, namely the unit cube $[0,1]^3$, and then turn to more complicated domains, which are shown in Figure 7.2. The first one is a thick quarter of ring; note that this solid has a trivial geometry on the third direction. The second one is the solid of revolution obtained by the 2D quarter of ring. Specifically, we performed a $\pi/2$ revolution around the axis having direction $(0,1,0)$ and passing through $(-1,-1,-1)$. We emphasize that here the geometry is nontrivial along all directions. In the cube case, we set $b = \texttt{randn}(n^3, 1)$ for computational ease, while in the other two cases b is the vector representing the function $f(x,y,z) = 2\left(x^2 - x\right) + 2\left(y^2 - y\right) + 2\left(z^2 - z\right)$. We again set K as the identity matrix in all cases.

We report in Table 7.6 the performances of FD and ADI (where the parameters are chosen as in [20]) relative to the cube domain. We can see that the computational time required by FD is independent of the degree p. In fact, the timings look impressive

	CG + $\mathcal{P}$ Iterations / Time (sec)			
h^{-1}	$p=2$	$p=3$	$p=4$	$p=5$
128	125 / 0.18	155 / 0.27	186 / 0.48	216 / 0.74
256	189 / 1.33	236 / 2.10	280 / 3.30	320 / 4.56
512	279 / 10.59	345 / 15.83	406 / 21.71	446 / 29.45
1024	404 / 99.15	487 / 140.64	556 / 174.55	587 / 203.66

	CG + $\mathcal{P}_J$ Iterations / Time (sec)			
h^{-1}	$p=2$	$p=3$	$p=4$	$p=5$
128	200 (5) / 0.95	260 (5) / 1.40	324 (5) / 2.16	380 (6) / 3.49
256	346 (6) / 7.34	429 (6) / 10.25	522 (6) / 14.91	620 (6) / 20.66
512	582 (6) / 49.93	714 (6) / 68.24	846 (6) / 92.16	870 (7) / 127.39
1024	771 (7) / 319.56	962 (7) / 438.11	1184 (7) / 605.72	1366 (7) / 834.23

	CG + IC Iterations / Time (sec)			
h^{-1}	$p=2$	$p=3$	$p=4$	$p=5$
128	92 / 0.15	55 / 0.17	45 / 0.24	38 / 0.33
256	180 / 1.27	114 / 1.37	90 / 1.70	73 / 2.08
512	354 / 10.62	253 / 12.13	193 / 13.64	151 / 15.40
1024	695 / 86.53	597 / 111.67	444 / 121.18	339 / 132.59

TABLE 7.5

Plate with hole domain. Performance of CG preconditioned by FD (upper table), by ADI (middle table) and by Incomplete Cholesky (lower table).

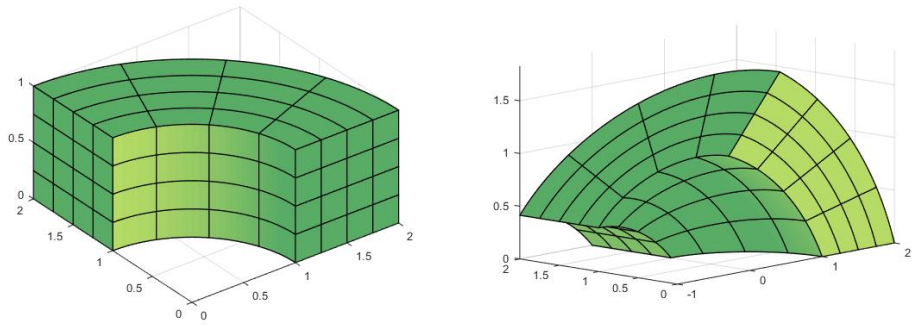

FIG. 7.2. *Thick ring and revolved ring domains*

and show the great efficiency of this approach. We emphasize that, on the finest discretization level, problems with more than one billion variables are solved in slightly more than five minutes, regardless of p. On the other hand, the ADI solver shows a considerably worse performance. Indeed, while the results confirm that this approach is robust w.r.t. h and p, the timings are always a couple of orders of magnitude greater than those obtained with FD. A comparison with Table 7.1 shows also that, as expected, the number of ADI iterations is higher than in the 2D case.

We mention that, for ADI, we also tested with the greedy choice of the parameters defined by (5.5)-(5.6). This choice yields about 10%-20% less iterations than the standard approach. Despite being an effective strategy, the improvement is not dramatic and the FD method is still much more efficient. In the following tests, we always consider the parameters from [20].

We now consider the problems with nontrivial geometries, where the two methods

	FD Direct Solver Time (sec)					
h^{-1}	$p=1$	$p=2$	$p=3$	$p=4$	$p=5$	$p=6$
128	0.19	0.15	0.17	0.17	0.18	0.17
256	1.80	1.87	2.05	1.90	2.10	1.82
512	23.02	22.58	23.45	23.26	23.65	21.89
1024	331.01	316.15	328.65	318.42	331.06	310.31

	3D ADI Iterations / Time (sec)					
h^{-1}	$p=1$	$p=2$	$p=3$	$p=4$	$p=5$	$p=6$
128	57/ 16.91	57/ 22.50	57/ 18.31	65/ 22.12	65/ 24.30	65/ 22.35
256	65/ 177.87	65/ 256.44	65/ 180.94	65/ 189.45	73/ 239.31	73/ 199.29
512	73/1872.01	73/2350.32	73/2239.23	73/2354.14	81/2714.25	81/1708.36

TABLE 7.6

Cube domain. Performance of FD (upper table) and of ADI (lower table). We did not run ADI on the finest level due to memory limitations.

	CG + $\mathcal{P}$ Iterations / Time (sec)				
h^{-1}	$p=2$	$p=3$	$p=4$	$p=5$	$p=6$
32	26 / 0.19	26 / 0.38	26 / 0.75	26 / 1.51	26 / 2.64
64	27 / 1.43	27 / 3.35	27 / 6.59	27 / 12.75	27 / 21.83
128	28 / 14.14	28 / 32.01	28 / 61.22	*	*

	CG + $\mathcal{P}_J$ Iterations / Time (sec)				
h^{-1}	$p=2$	$p=3$	$p=4$	$p=5$	$p=6$
32	26 (7) / 0.88	26 (7) / 1.20	26 (7) / 1.71	26 (7) / 2.62	27 (8) / 4.08
64	27 (7) / 7.20	27 (8) / 10.98	27 (8) / 14.89	27 (8) / 21.81	27 (8) / 30.56
128	28 (8) / 99.01	28 (8) / 98.39	28 (8) / 143.45	*	*

	CG + IC Iterations / Time (sec)				
h^{-1}	$p=2$	$p=3$	$p=4$	$p=5$	$p=6$
32	21 / 0.37	15 / 1.17	12 / 3.41	10 / 9.43	9 / 24.05
64	37 / 4.26	28 / 13.23	22 / 33.96	18 / 88.94	16 / 215.31
128	73 / 65.03	51 / 163.48	41 / 385.54	*	*

TABLE 7.7

Thick quarter of ring domain. Performance of CG preconditioned by the direct method (upper table), by ADI (middle table) and by Incomplete Cholesky (lower table).

are used as preconditioners for CG. In the case of ADI, we again set $\epsilon = 10^{-1}$ for both problems. As in the 2D case, we also consider a standard Incomplete Cholesky (IC) preconditioner (no reordering is used in this case, as the resulting performance is better than when using the standard reorderings available in MATLAB).

In Table 7.7 we report the results for the thick quarter ring while in Table 7.8 we report the results for the revolved ring. The symbol "*" denotes the cases in which even assembling the system matrix $\mathcal{A}$ was unfeasible due to memory limitations. From these results, we infer that most of the conclusions drawn for the 2D case still hold in 3D. In particular, both Sylvester-based preconditioners yield a better performance than the IC preconditioner, especially for small h.

Somewhat surprisingly, however, the CPU times show a stronger dependence on p than in the 2D case, and the performance gap between the ADI and the FD approach is not as large as for the cube domain. This is due to the cost of the residual computation in the CG iteration (a sparse matrix-vector product, costing $O(p^3n^3)$ FLOPs). This step represents now a significant computational effort in the

	CG + $\mathcal{P}$ Iterations / Time (sec)				
h^{-1}	$p=2$	$p=3$	$p=4$	$p=5$	$p=6$
32	40 / 0.27	41 / 0.63	41 / 1.24	42 / 2.38	42 / 4.13
64	44 / 2.30	44 / 5.09	45 / 10.75	45 / 20.69	45 / 35.11
128	47 / 23.26	47 / 55.34	47 / 101.94	*	*

	CG + $\mathcal{P}_J$ Iterations / Time (sec)				
h^{-1}	$p=2$	$p=3$	$p=4$	$p=5$	$p=6$
32	40 (7) / 1.39	41 (7) / 1.93	41 (7) / 2.67	42 (7) / 4.17	42 (8) / 6.25
64	44 (7) / 11.82	44 (8) / 16.96	45 (8) / 24.31	45 (8) / 35.76	45 (8) / 49.89
128	47 (8) / 170.69	47 (8) / 168.45	47 (9) / 239.07	*	*

	CG + IC Iterations / Time (sec)				
h^{-1}	$p=2$	$p=3$	$p=4$	$p=5$	$p=6$
32	24 / 0.44	18 / 1.28	15 / 3.61	12 / 9.63	11 / 24.57
64	47 / 5.19	35 / 14.95	28 / 37.33	24 / 94.08	20 / 222.09
128	94 / 81.65	71 / 211.53	57 / 464.84	*	*

TABLE 7.8

Revolved quarter of ring domain. Performance of CG preconditioned by the direct method (upper table), by ADI (middle table) and by Incomplete Cholesky (lower table).

h^{-1}	$p=2$	$p=3$	$p=4$	$p=5$	$p=6$
32	25.60	13.34	7.40	4.16	2.44
64	22.69	11.26	5.84	3.32	1.88
128	25.64	13.09	6.92	*	*

TABLE 7.9

Percentage of time spent in the application of the 3D FD preconditioner with respect to the overall CG time. Revolved ring domain.

overall CG performance. In fact, our numerical experience shows that the 3D FD method is so efficient that the time spent in the preconditioning step is often negligible w.r.t. the time required for the residual computation. This effect is clearly shown in Table 7.9, where we we report the percentage of time spent in the application of the preconditioner when compared with the overall time of CG, in the case of the revolved ring domain. Interestingly, this percentage is almost constant w.r.t. h up to the finest discretization level, corresponding to about 2 million degrees-of-freedom.

7.3. Multi-patch experiments. In this section we consider a multi-patch 2D problem. We consider the L-shaped domain shown in Figure 7.3, discretized with 3 patches and imposing C^0 continuity at the interfaces. We solve this problem using the additive overlapping Schwarz preconditioner described in Section 6. Here Ω is split into two rectangular subdomains, which overlap on Ω_2. Despite its apparent simplicity, this problem contains all the relevant ingredients to test the validity of our approach in the multi-patch case. Note in particular that the Jacobian of the geometry mapping on the subdomains is not the identity matrix, due to stretching in either direction. This stretching could be easily incorporated into the preconditioner, but we avoid doing that in order to mimic the effect of nontrivial geometry.

We consider both the exact preconditioner, where the local systems are solved with MATLAB's direct solver, and the inexact one (6.2), with $\widetilde{\mathcal{A}}_i = \mathcal{P}_i$, implemented by the FD solver. For completeness, we also consider the IC preconditioner, like in previous experiments. The results are shown in Table 7.10. At the coarsest level and for low degree, the three approaches perform similarly, in terms of CPU time. However, as

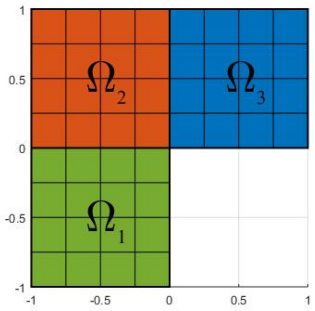

FIG. 7.3. *L-shaped domain domain (3 patches)*

	CG + $\mathcal{P}_{EAS}$ Iterations / Time (sec)				
h^{-1}	$p=1$	$p=2$	$p=3$	$p=4$	$p=5$
128	3 / 0.88	2 / 1.17	2 / 2.51	2 / 3.82	2 / 6.14
256	2 / 2.98	2 / 6.43	2 / 16.44	2 / 24.98	2 / 41.27
512	2 / 17.42	2 / 34.83	2 / 128.71	2 / 168.33	2 / 334.95
1024	2 / 110.55	2 / 266.03	2 / 998.72	2 / 2028.72	2 / 2530.75

	CG + $\mathcal{P}_{IAS}$ Iterations / Time (sec)				
h^{-1}	$p=1$	$p=2$	$p=3$	$p=4$	$p=5$
128	20 / 0.79	20 / 0.54	20 / 0.64	19 / 0.74	19 / 0.86
256	19 / 1.46	20 / 1.69	19 / 1.91	19 / 2.20	19 / 2.55
512	19 / 7.32	19 / 7.50	19 / 8.11	19 / 9.07	19 / 10.23
1024	19 / 44.40	19 / 44.57	19 / 49.50	18 / 61.42	18 / 55.84

	CG + IC Iterations / Time (sec)				
h^{-1}	$p=1$	$p=2$	$p=3$	$p=4$	$p=5$
128	144 / 0.58	94 / 0.57	69 / 0.58	56 / 0.71	46 / 0.89
256	280 / 3.76	180 / 4.07	132 / 4.65	106 / 5.46	87 / 6.23
512	544 / 31.07	348 / 31.86	254 / 36.10	203 / 42.44	166 / 48.50
1024	1052 / 237.01	673 / 246.65	491 / 273.97	392 / 325.38	321 / 470.61

TABLE 7.10

L-shaped domain. Performance of CG preconditioned by $\mathcal{P}_{EAS}$ (upper table), by $\mathcal{P}_{IAS}$ (middle table) and by Incomplete Cholesky (lower table).

expected the $\mathcal{P}_{IAS}$ preconditioner scales much better than the others w.r.t. h and p. In particular, the number of iterations is independent of these parameters.

We also tested with $\mathcal{P}_{IAS}$, when however the matrices $\widetilde{\mathcal{A}}_i^{-1}$ represent a fixed number of preconditioned CG iterations. The results are comparable with the ones obtained for $\widetilde{\mathcal{A}}_i = \mathcal{P}_i$, although slightly worse, and hence they are not shown.

8. Conclusions. In this work we have analyzed and tested the use of fast solvers for Sylvester-like equations as preconditioners for isogeometric discretizations.

We considered here a Poisson problem on a single-patch domain, and we focused on the k-method, i.e., splines with maximal smoothness. The considered preconditioner $\mathcal{P}$ is robust w.r.t. h and p, and we have compared two popular methods for its application. We found that the FD direct solver, especially in 3D, is by far more effective than the ADI iterative solver. Both approaches easily outperform a simple-minded Incomplete Cholesky preconditioner.

Our conclusion is that the use of the FD method is, likely, the best possible choice to compute the action of the operator $\mathcal{P}^{-1}$. In fact, even if a more efficient solver was available, its employment would not necessary yield a relevant improvement in the overall performance, since in our experiments the cost to apply the FD solver is already negligible w.r.t. the cost of a single matrix-vector product. This is, then, a very promising preconditioning stage in an iterative solver for isogeometric discretizations. As we showed, this preconditioner can be easily combined with a domain decomposition strategy to solve multi-patch problems with conforming discretization. In a forthcoming paper we will further study the role of the geometry parametrization on the performance of the approaches based on Sylvester equation solvers, and propose possible strategies to improve it.

Acknowledgments. The authors would like to thank Valeria Simoncini for fruitful discussions on the topics of the paper. The authors were partially supported by the European Research Council through the FP7 Ideas Consolodator Grant *HIGEOM* n.616563. This support is gratefully acknowledged.

REFERENCES

[1] J. Ballani and L. Grasedyck, *A projection method to solve linear systems in tensor format*, Numerical Linear Algebra with Applications, 20 (2013), pp. 27–43.

[2] R. H. Bartels and G. W. Stewart, *Solution of the matrix equation $AX+ XB= C$*, Communications of the ACM, 15 (1972), pp. 820–826.

[3] Y. Bazilevs, C. Michler, V. M. Calo, and T. J. R. Hughes, *Isogeometric variational multiscale modeling of wall-bounded turbulent flows with weakly enforced boundary conditions on unstretched meshes*, Computer Methods in Applied Mechanics and Engineering, 199 (2010), pp. 780–790.

[4] L. Beirão da Veiga, A. Buffa, G. Sangalli, and R. Vázquez, *Mathematical analysis of variational isogeometric methods*, Acta Numerica, 23 (2014), pp. 157–287.

[5] L. Beirão da Veiga, D. Cho, L. F. Pavarino, and S. Scacchi, *Overlapping Schwarz methods for isogeometric analysis*, SIAM Journal on Numerical Analysis, 50 (2012), pp. 1394–1416.

[6] ———, *BDDC preconditioners for isogeometric analysis*, Mathematical Models and Methods in Applied Sciences, 23 (2013), pp. 1099–1142.

[7] P. Benner and T. Breiten, *Low rank methods for a class of generalized Lyapunov equations and related issues*, Numerische Mathematik, 124 (2013), pp. 441–470.

[8] P. Benner, R. C. Li, and N. Truhar, *On the ADI method for Sylvester equations*, Journal of Computational and Applied Mathematics, 233 (2009), pp. 1035–1045.

[9] M. Bercovier and I. Soloveichik, *Overlapping non matching meshes domain decomposition method in isogeometric analysis*, arXiv preprint arXiv:1502.03756, (2015).

[10] A. Buffa, H. Harbrecht, A. Kunoth, and G. Sangalli, *BPX-preconditioning for isogeometric analysis*, Computer Methods in Applied Mechanics and Engineering, 265 (2013), pp. 63–70.

[11] N. Collier, L. Dalcin, D. Pardo, and V. M. Calo, *The cost of continuity: performance of iterative solvers on isogeometric finite elements*, SIAM Journal on Scientific Computing, 35 (2013), pp. A767–A784.

[12] N. Collier, D. Pardo, L. Dalcin, M. Paszynski, and V. M. Calo, *The cost of continuity: a study of the performance of isogeometric finite elements using direct solvers*, Computer Methods in Applied Mechanics and Engineering, 213 (2012), pp. 353–361.

[13] J. A. Cottrell, T. J. R. Hughes, and Y. Bazilevs, *Isogeometric analysis: toward integration of CAD and FEA*, John Wiley & Sons, 2009.

[14] J. A. Cottrell, T. J. R. Hughes, and A. Reali, *Studies of refinement and continuity in isogeometric structural analysis*, Computer Methods in Applied Mechanics and Engineering, 196 (2007), pp. 4160–4183.

[15] T. Damm, *Direct methods and ADI-preconditioned Krylov subspace methods for generalized Lyapunov equations*, Numerical Linear Algebra with Applications, 15 (2008), pp. 853–871.

[16] C. De Falco, A. Reali, and R. Vázquez, *Geopdes: a research tool for isogeometric analysis of pdes*, Advances in Engineering Software, 42 (2011), pp. 1020–1034.

[17] J. W. Demmel, *Applied numerical linear algebra*, SIAM, 1997.
[18] M. O. Deville, P. F. Fischer, and E. H. Mund, *High-order methods for incompressible fluid flow*, Cambridge University Press, 2002.
[19] M. Donatelli, C. Garoni, C. Manni, S. Serra-Capizzano, and H. Speleers, *Robust and optimal multi-iterative techniques for IgA Galerkin linear systems*, Computer Methods in Applied Mechanics and Engineering, 284 (2015), pp. 230–264.
[20] J. Douglas, *Alternating direction methods for three space variables*, Numerische Mathematik, 4 (1962), pp. 41–63.
[21] N. S. Ellner and E. L. Wachspress, *Alternating direction implicit iteration for systems with complex spectra*, SIAM journal on numerical analysis, 28 (1991), pp. 859–870.
[22] K. P. S. Gahalaut, J. K. Kraus, and S. K. Tomar, *Multigrid methods for isogeometric discretization*, Computer Methods in Applied Mechanics and Engineering, 253 (2013), pp. 413–425.
[23] L. Gao, *Kronecker Products on Preconditioning*, PhD thesis, King Abdullah University of Science and Technology, 2013.
[24] G. H. Golub and C. F. Van Loan, *Matrix computations*, Johns Hopkins University Press, 2012.
[25] L. Grasedyck, *Existence and computation of low Kronecker-rank approximations for large linear systems of tensor product structure*, Computing, 72 (2004), pp. 247–265.
[26] C. Hofreither, S. Takacs, and W. Zulehner, *A robust multigrid method for isogeometric analysis using boundary correction*, Tech. Report 33, NFN, 2015.
[27] T. J. R. Hughes, J. A. Cottrell, and Y. Bazilevs, *Isogeometric analysis: CAD, finite elements, NURBS, exact geometry and mesh refinement*, Computer Methods in Applied Mechanics and Engineering, 194 (2005), pp. 4135–4195.
[28] S. K. Kleiss, C. Pechstein, B. Jüttler, and S. Tomar, *IETI–Isogeometric tearing and interconnecting*, Computer Methods in Applied Mechanics and Engineering, 247 (2012), pp. 201–215.
[29] D. Kressner, M. Plešinger, and C. Tobler, *A preconditioned low-rank CG method for parameter-dependent Lyapunov matrix equations*, Numerical Linear Algebra with Applications, 21 (2014), pp. 666–684.
[30] D. Kressner and C. Tobler, *Krylov subspace methods for linear systems with tensor product structure*, SIAM Journal on Matrix Analysis and Applications, 31 (2010), pp. 1688–1714.
[31] B. W. Li, S. Tian, Y. S. Sun, and Z. M. Hu, *Schur-decomposition for 3D matrix equations and its application in solving radiative discrete ordinates equations discretized by Chebyshev collocation spectral method*, Journal of Computational Physics, 229 (2010), pp. 1198–1212.
[32] A. Lu and E. L. Wachspress, *Solution of Lyapunov equations by alternating direction implicit iteration*, Computers & Mathematics with Applications, 21 (1991), pp. 43–58.
[33] R. E. Lynch, J. R: Rice, and D. H. Thomas, *Direct solution of partial difference equations by tensor product methods*, Numerische Mathematik, 6 (1964), pp. 185–199.
[34] D. W. Peaceman and H. H. Rachford, *The numerical solution of parabolic and elliptic differential equations*, Journal of the Society for Industrial & Applied Mathematics, 3 (1955), pp. 28–41.
[35] T. Penzl, *A cyclic low-rank Smith method for large sparse Lyapunov equations*, SIAM Journal on Scientific Computing, 21 (2000), pp. 1401–1418.
[36] V. Simoncini, *Computational methods for linear matrix equations*, to appear on SIAM Review, (2013).
[37] L. Sorber, M. Van Barel, and L. De Lathauwer, *Tensorlab v2. 0*, Available online, URL: www. tensorlab. net, (2014).
[38] S. Takacs and T. Takacs, *Approximation error estimates and inverse inequalities for B-splines of maximum smoothness*, arXiv preprint arXiv:1502.03733, (2015).
[39] E. L. Wachspress, *Optimum alternating-direction-implicit iteration parameters for a model problem*, Journal of the Society for Industrial & Applied Mathematics, 10 (1962), pp. 339–350.
[40] ———, *Extended application of alternating direction implicit iteration model problem theory*, Journal of the Society for Industrial & Applied Mathematics, 11 (1963), pp. 994–1016.
[41] ———, *Three-variable alternating-direction-implicit iteration*, Computers & Mathematics with Applications, 27 (1994), pp. 1–7.
[42] ———, *The ADI model problem*, Springer, 2013.